arXiv:1702.00146v1 [cs.CG] 1 Feb 2017

Untangling Planar Curves*

Hsien-Chih Chang Jeff Erickson

Department of Computer Science
University of Illinois, Urbana-Champaign
{hchang17, jeffe}@illinois.edu

Submitted to *Discrete & Computational Geometry* — July 16, 2016
Revised and resubmitted — February 2, 2017

Abstract

Any generic closed curve in the plane can be transformed into a simple closed curve by a finite sequence of local transformations called *homotopy moves*. We prove that simplifying a planar closed curve with n self-crossings requires $\Theta(n^{3/2})$ homotopy moves in the worst case. Our algorithm improves the best previous upper bound $O(n^2)$, which is already implicit in the classical work of Steinitz; the matching lower bound follows from the construction of closed curves with large *defect*, a topological invariant of generic closed curves introduced by Aicardi and Arnold. Our lower bound also implies that $\Omega(n^{3/2})$ facial electrical transformations are required to reduce any plane graph with treewidth $\Omega(\sqrt{n})$ to a single vertex, matching known upper bounds for rectangular and cylindrical grid graphs. More generally, we prove that transforming one immersion of k circles with at most n self-crossings into another requires $\Theta(n^{3/2} + nk + k^2)$ homotopy moves in the worst case. Finally, we prove that transforming one noncontractible closed curve to another on any orientable surface requires $\Omega(n^2)$ homotopy moves in the worst case; this lower bound is tight if the curve is homotopic to a simple closed curve.

"It's hardly fair," muttered Hugh, "to give us such a jumble as this to work out!"

— Lewis Carroll, *A Tangled Tale*, Knot X (1885)

*Work on this paper was partially supported by NSF grant CCF-1408763. A preliminary version of this paper was presented at the 32nd International Symposium on Computational Geometry [12]. See http://jeffe.cs.illinois.edu/pubs/tangle.pdf for the most recent version of this paper.

1 Introduction

Any generic closed curve in the plane can be transformed into a simple closed curve by a finite sequence of the following local operations:

- **1→0**: Remove an empty loop.
- **2→0**: Separate two subpaths that bound an empty bigon.
- **3→3**: Flip an empty triangle by moving one subpath over the opposite intersection point.

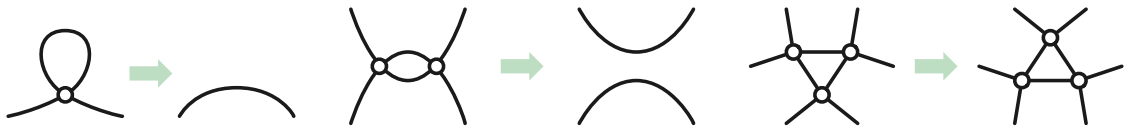

Figure 1.1. Homotopy moves 1→0, 2→0, and 3→3.

See Figure 1.1. Each of these operations can be performed by continuously deforming the curve within a small neighborhood of one face; consequently, we call these operations and their inverses ***homotopy moves***. Our notation is nonstandard but mnemonic; the numbers before and after each arrow indicate the number of local vertices before and after the move. Homotopy moves are "shadows" of the classical Reidemeister moves used to manipulate knot and link diagrams [4, 50].

We prove that $\Theta(n^{3/2})$ homotopy moves are sometimes necessary and always sufficient to simplify a closed curve in the plane with n self-crossings. Before describing our results in more detail, we review several previous results.

1.1 Past Results

An algorithm to simplify any planar closed curve using at most $O(n^2)$ homotopy moves is implicit in Steinitz's proof that every 3-connected planar graph is the 1-skeleton of a convex polyhedron [55, 56]. Specifically, Steinitz proved that any non-simple closed curve (in fact, any 4-regular plane graph) with no empty loops contains a *bigon* ("Spindel"): a disk bounded by a pair of simple subpaths that cross exactly twice, where the endpoints of the (slightly extended) subpaths lie outside the disk. Steinitz then proved that any *minimal* bigon ("irreduzible Spindel") can be transformed into an empty bigon using a sequence of 3→3 moves, each removing one triangular face from the bigon, as shown in Figure 1.2. Once the bigon is empty, it can be deleted with a single 2→0 move. See Grünbaum [30], Hass and Scott [34], Colin de Verdière *et al.* [17], or Nowik [47] for more modern treatments of Steinitz's technique. The $O(n^2)$ upper bound also follows from algorithms for *regular* homotopy, which forbids 0↔1 moves, by Francis [25], Vegter [61] (for polygonal curves), and Nowik [47].

The $O(n^2)$ upper bound can also be derived from an algorithm of Feo and Provan [24] for reducing a *plane* graph to a single edge by *electrical transformations*: degree-1 reductions, series-parallel reductions, and ΔY-transformations. (We consider electrical transformations in more detail in Section 3.) Any curve divides the plane into regions, called its *faces*. The *depth* of a face is its distance to the outer face in the dual graph of the curve. Call a homotopy move *positive* if it decreases the sum of the face depths; in particular, every 1→0 and 2→0 move is positive. A key technical lemma of Feo and Provan implies that every non-simple curve in the plane admits a positive homotopy move [24, Theorem 1]. Thus, the sum of the face depths is an upper bound on the minimum number of moves required to simplify the curve. Euler's formula implies that every curve with n crossings has $O(n)$ faces, and each of these faces has depth $O(n)$.

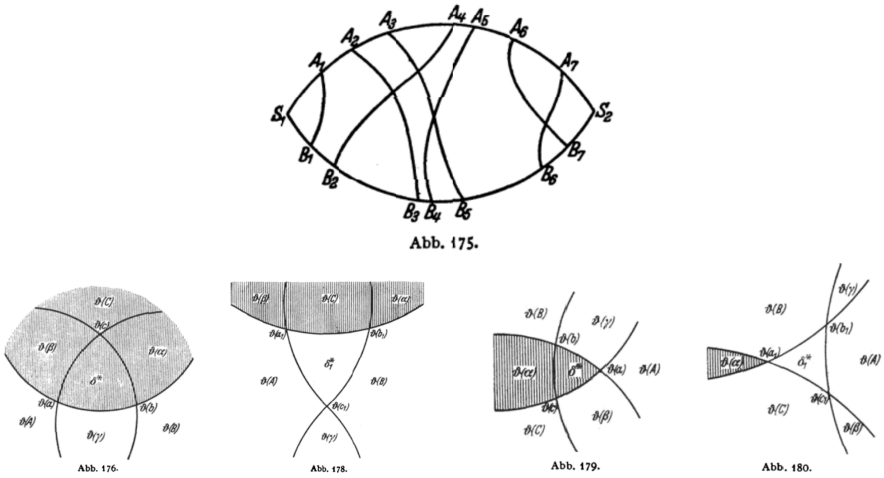

Figure 1.2. Top: A minimal bigon. Bottom: 3→3 moves removing triangles from the side or the end of a (shaded) minimal bigon. All figures are from Steinitz and Rademacher [56].

Gitler [27] conjectured that a variant of Feo and Provan's algorithm that always makes the *deepest* positive move requires only $O(n^{3/2})$ moves. Song [54] observed that if Feo and Provan's algorithm always chooses the *shallowest* positive move, it can be forced to make $\Omega(n^2)$ moves even when the input curve can be simplified using only $O(n)$ moves.

Tight bounds are known for two special cases where some homotopy moves are forbidden. First, Nowik [47] proved a tight $\Omega(n^2)$ lower bound for regular homotopy. Second, Khovanov [41] defined two curves to be *doodle equivalent* if one can be transformed into the other using $1\leftrightarrow 0$ and $2\leftrightarrow 0$ moves. Khovanov [41] and Ito and Takimura [39] independently proved that any planar curve can be transformed into its unique equivalent doodle with the smallest number of vertices, using only $1\to 0$ and $2\to 0$ moves. Thus, two doodle equivalent curves are connected by a sequence of $O(n)$ moves, which is obviously tight.

Looser bounds are also known for the minimum number of Reidemeister moves needed to reduce a diagram of the unknot [32, 42], to separate the components of a split link [36], or to move between two equivalent knot diagrams [19, 35].

1.2 New Results

In Section 2, we derive an $\Omega(n^{3/2})$ lower bound using a numerical curve invariant called *defect*, introduced by Arnold [8, 9] and Aicardi [1]. Each homotopy move changes the defect of a closed curve by at most 2. The lower bound therefore follows from constructions of Hayashi *et al.* [35, 37] and Even-Zohar *et al.* [22] of closed curves with defect $\Omega(n^{3/2})$. We simplify and generalize their results by computing the defect of the standard planar projection of any $p \times q$ torus knot where either $p \bmod q = 1$ or $q \bmod p = 1$. Our calculations imply that for any integer p, reducing the standard projection of the $p \times (p+1)$ torus knot requires at least $\binom{p+1}{3} \geq n^{3/2}/6 - O(n)$ homotopy moves. Finally, using winding-number arguments, we prove that in the worst case, simplifying an arrangement of k closed curves requires $\Omega(n^{3/2} + nk)$ homotopy moves, with an additional $\Omega(k^2)$ term if the target configuration is specified in advance.

In Section 3, we provide a proof, based on arguments of Truemper [58] and Noble and Welsh [46], that reducing a *unicursal* plane graph G—one whose medial graph is the image of a single closed curve—using facial electrical transformations requires at least as many steps as reducing the medial graph of G to a simple closed curve using homotopy moves. The homotopy lower bound from Section 2 then implies that reducing any n-vertex plane graph with treewidth $\Omega(\sqrt{n})$ requires $\Omega(n^{3/2})$ facial electrical

transformations. This lower bound matches known upper bounds for rectangular and cylindrical grid graphs.

We develop a new algorithm to simplify any closed curve in $O(n^{3/2})$ homotopy moves in Section 4. First we describe an algorithm that uses $O(D)$ moves, where D is the sum of the face depths of the input curve. At a high level, our algorithm can be viewed as a variant of Steinitz's algorithm that empties and removes *loops* instead of bigons. We then extend our algorithm to *tangles*: collections of boundary-to-boundary paths in a closed disk. Our algorithm simplifies a tangle as much as possible in $O(D + ns)$ moves, where D is the sum of the depths of the tangle's faces, s is the number of paths, and n is the number of intersection points. Then, we prove that for any curve with maximum face depth $\Omega(\sqrt{n})$, we can find a simple closed curve whose interior tangle has m interior vertices, at most $\sqrt{m}$ paths, and maximum face depth $O(\sqrt{n})$. Simplifying this tangle and then recursively simplifying the resulting curve requires a total of $O(n^{3/2})$ moves. We show that this simplifying sequence of homotopy moves can be computed in $O(1)$ amortized time per move, assuming the curve is presented in an appropriate graph data structure. We conclude this section by proving that any arrangement of k closed curves can be simplified in $O(n^{3/2} + nk)$ homotopy moves, or in $O(n^{3/2} + nk + k^2)$ homotopy moves if the target configuration is specified in advance, precisely matching our lower bounds for all values of n and k.

Finally, in Section 5, we consider curves on surfaces of higher genus. We prove that $\Omega(n^2)$ homotopy moves are required in the worst case to transform one non-contractible closed curve to another on the torus, and therefore on any orientable surface. Results of Hass and Scott [33] imply that this lower bound is tight if the non-contractible closed curve is homotopic to a simple closed curve.

1.3 Definitions

A ***closed curve*** in a surface M is a continuous map $\gamma\colon S^1 \to M$. In this paper, we consider only *generic* closed curves, which are injective except at a finite number of self-intersections, each of which is a transverse double point; closed curves satisfying these conditions are called *immersions* of the circle. A closed curve is ***simple*** if it is injective. For most of the paper, we consider only closed curves in the plane; we consider more general surfaces in Section 5.

The image of any non-simple closed curve has a natural structure as a 4-regular plane graph. Thus, we refer to the self-intersection points of a curve as its ***vertices***, the maximal subpaths between vertices as ***edges***, and the components of the complement of the curve as its ***faces***. Two curves γ and γ' are *isomorphic* if their images define combinatorially equivalent maps; we will not distinguish between isomorphic curves.

A *corner* of γ is the intersection of a face of γ and a small neighborhood of a vertex of γ. A ***loop*** in a closed curve γ is a subpath of γ that begins and ends at some vertex x, intersects itself only at x, and encloses exactly one corner at x. A ***bigon*** in γ consists of two simple interior-disjoint subpaths of γ with the same endpoints and enclose one corner at each of those endpoints. A loop or bigon is ***empty*** if its interior does not intersect γ. Notice that a 1→0 move is applied to an empty loop, and a 2→0 move is applied on an empty bigon.

We adopt a standard sign convention for vertices first used by Gauss [26]. Choose an arbitrary basepoint $\gamma(0)$ and orientation for the curve. For each vertex x, we define $\operatorname{sgn}(x) = +1$ if the first traversal through the vertex crosses the second traversal from right to left, and $\operatorname{sgn}(x) = -1$ otherwise. See Figure 1.3.

A ***homotopy*** between two curves γ and γ' in surface M is a continuous function $H\colon S^1 \times [0,1] \to M$ such that $H(\cdot, 0) = \gamma$ and $H(\cdot, 1) = \gamma'$. Any homotopy H describes a continuous deformation from γ to γ', where the second argument of H is "time". Each homotopy move can be executed by a homotopy. Conversely, Alexander's simplicial approximation theorem [3], together with combinatorial arguments of Alexander and Briggs [4] and Reidemeister [50], imply that any generic homotopy between two closed

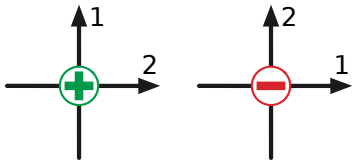

Figure 1.3. Gauss's sign convention.

curves can be decomposed into a finite sequence of homotopy moves. Two curves are *homotopic*, or in the same *homotopy class*, if there is a homotopy from one to the other. All closed curves in the plane are homotopic.

A ***multicurve*** is an immersion of one or more disjoint circles; in particular, a ***k-curve*** is an immersion of k disjoint circles. A multicurve is *simple* if it is injective, or equivalently, if it can be decomposed into pairwise disjoint simple closed curves. The image of any multicurve in the plane is the disjoint union of simple closed curves and 4-regular plane graphs. A ***component*** of a multicurve γ is any multicurve whose image is a connected component of the image of γ. We call the individual closed curves that comprise a multicurve its ***constituent curves***; see Figure 1.4. The definition of homotopy and the decomposition of homotopies into homotopy moves extend naturally to multicurves.

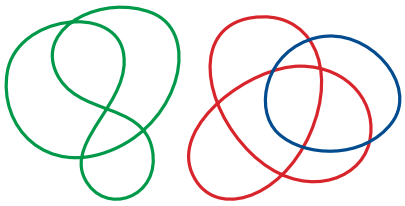

Figure 1.4. A multicurve with two components and three constituent curves, one of which is simple.

2 Lower Bounds

2.1 Defect

To prove our main lower bound, we consider a numerical invariant of closed curves in the plane introduced by Arnold [8, 9] and Aicardi [1] called ***defect***. Polyak [49] proved that defect can be computed—or for our purposes, defined—as follows:

$$\mathit{defect}(\gamma) := -2 \sum_{x \between y} \operatorname{sgn}(x) \cdot \operatorname{sgn}(y).$$

Here the sum is taken over all *interleaved* pairs of vertices of γ: two vertices $x \neq y$ are interleaved, denoted $\boldsymbol{x \between y}$, if they alternate in cyclic order—x, y, x, y—along γ. (The factor of -2 is a historical artifact, which we retain only to be consistent with Arnold's original definitions [8, 9].) Even though the signs of individual vertices depend on the basepoint and orientation of the curve, the defect of a curve is independent of those choices. Moreover, the defect of any curve is preserved by any homeomorphism from the plane (or the sphere) to itself, including reflection.

Trivially, every simple closed curve has defect zero. Straightforward case analysis [49] implies that any single homotopy move changes the defect of a curve by at most 2; the various cases are listed below and illustrated in Figure 2.1.

- A 1→0 move leaves the defect unchanged.
- A 2→0 move decreases the defect by 2 if the two disappearing vertices are interleaved, and leaves the defect unchanged otherwise.
- A 3→3 move increases the defect by 2 if the three vertices before the move contain an even number of interleaved pairs, and decreases the defect by 2 otherwise.

In light of this case analysis, the following lemma is trivial:

Lemma 2.1. *Simplifying any closed curve γ in the plane requires at least $|\mathit{defect}(\gamma)|/2$ homotopy moves.*

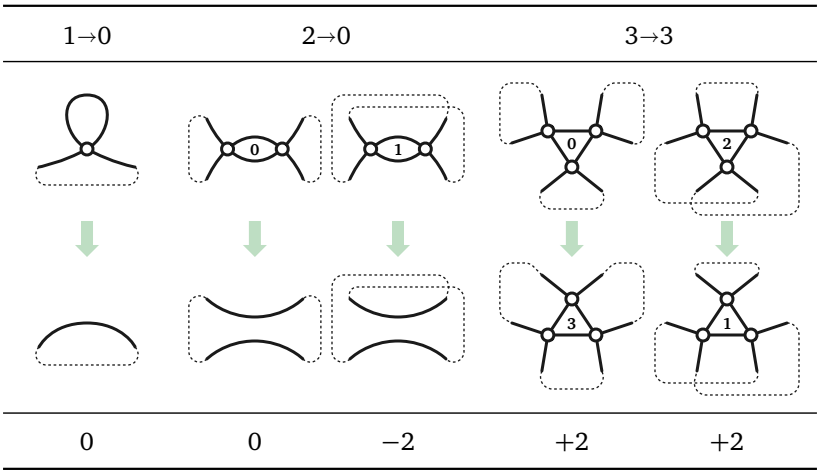

Figure 2.1. Changes to defect incurred by homotopy moves. Numbers in each figure indicate how many pairs of vertices are interleaved; dashed lines indicate how the rest of the curve connects.

2.2 Flat Torus Knots

For any relatively prime positive integers p and q, let $\boldsymbol{T(p,q)}$ denote the curve with the following parametrization, where θ runs from 0 to 2π:

$$T(p,q)(\theta) := \big((\cos(q\theta)+2)\cos(p\theta),\ (\cos(q\theta)+2)\sin(p\theta)\big)\,.$$

The curve $T(p,q)$ winds around the origin p times, oscillates q times between two concentric circles, and crosses itself exactly $(p-1)q$ times. We call these curves ***flat torus knots***.

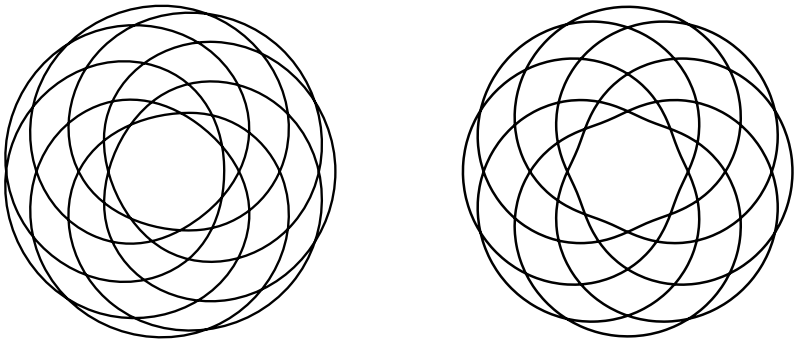

Figure 2.2. The flat torus knots $T(8,7)$ and $T(7,8)$.

Hayashi *et al.* [37, Proposition 3.1] proved that for any integer q, the flat torus knot $T(q+1,q)$ has defect $-2\binom{q}{3}$. Even-Zohar *et al.* [22] used a star-polygon representation of the curve $T(p,2p+1)$ as the basis for a universal model of random knots; in our notation, they proved that $\mathit{defect}(T(p,2p+1)) = 4\binom{p+1}{3}$ for any integer p. In this section we simplify and generalize both of these results to all flat torus knots $T(p,q)$ where either $q \bmod p = 1$ or $p \bmod q = 1$. For purposes of illustration, we cut $T(p,q)$ along a spiral path parallel to a portion of the curve, and then deform the p resulting subpaths, which we call *strands*, into a "flat braid" between two fixed diagonal lines. See Figure 2.3.

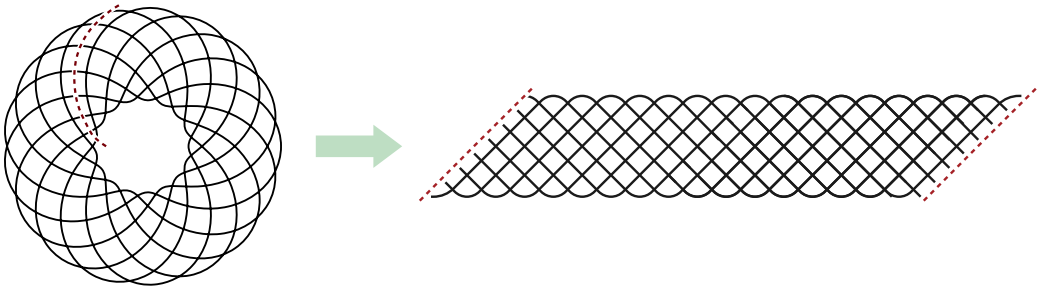

Figure 2.3. Transforming $T(8,17)$ into a flat braid.

Lemma 2.2. *$\mathit{defect}(T(p,ap+1)) = 2a\binom{p+1}{3}$ for all integers $a \geq 0$ and $p \geq 1$.*

Proof: The curve $T(p,1)$ can be reduced to a simple closed curve using only 1→0 moves, so its defect is zero. For the rest of the proof, assume $a \geq 1$.

We define a *stripe* of $T(p,ap+1)$ to be a subpath from some innermost point to the next outermost point, or equivalently, a subpath of any strand from the bottom to the top in the flat braid representation. Each stripe contains exactly $p-1$ crossings. A *block* of $T(p,ap+1)$ consists of $p(p-1)$ crossings in p consecutive stripes; within any block, each pair of strands intersects exactly twice. We can reduce $T(p,ap+1)$ to $T(p,(a-1)p+1)$ by straightening any block one strand at a time. Straightening the bottom strand of the block requires the following $\binom{p}{2}$ moves, as shown in Figure 2.4.

- $\binom{p-1}{2}$ 3→3 moves pull the bottom strand downward over one intersection point of every other pair of strands. Just before each 3→3 move, exactly one of the three pairs of the three relevant vertices is interleaved, so each move decreases the defect by 2.
- $(p-1)$ 2→0 moves eliminate a pair of intersection points between the bottom strand and every other strand. Each of these moves also decreases the defect by 2.

Altogether, straightening one strand decreases the defect by $2\binom{p}{2}$. Proceeding similarly with the other strands, we conclude that $\mathit{defect}(T(p,ap+1)) = \mathit{defect}(T(p,(a-1)p+1)) + 2\binom{p+1}{3}$. The lemma follows immediately by induction. □

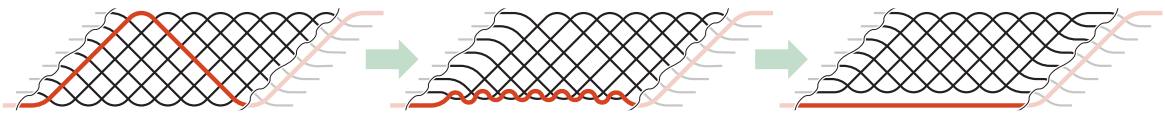

Figure 2.4. Straightening one strand in a block of $T(8,8a+1)$.

Lemma 2.3. *defect*$(T(aq+1,q)) = -2a\binom{q}{3}$ *for all integers* $a \geq 0$ *and* $q \geq 1$.

Proof: The curve $T(1,q)$ is simple, so its defect is trivially zero. For any positive integer a, we can transform $T(aq+1,q)$ into $T((a-1)q+1,q)$ by incrementally removing the innermost q *loops*. We can remove the first loop using $\binom{q}{2}$ homotopy moves, as shown in Figure 2.5. (The first transition in Figure 2.5 just reconnects the top left and top right endpoints of the flat braid.)

- $\binom{q-1}{2}$ 3→3 moves pull the left side of the loop to the right, over the crossings inside the loop. Just before each 3→3 move, the three relevant vertices contain two interleaved pairs, so each move *increases* the defect by 2.
- $(q-1)$ 2→0 moves pull the loop over $q-1$ strands. The strands involved in each move are oriented in opposite directions, so these moves leave the defect unchanged.
- Finally, we can remove the loop with a single 1→0 move, which does not change the defect.

Altogether, removing one loop increases the defect by $2\binom{q-1}{2}$. Proceeding similarly with the other loops, we conclude that *defect*$(T(aq+1,q)) =$ *defect*$(T((a-1)q+1,q)) - 2\binom{q}{3}$. The lemma follows immediately by induction. □

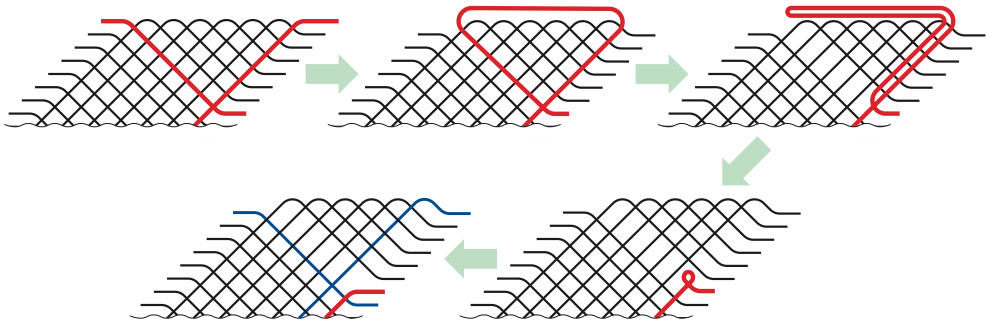

Figure 2.5. Removing one loop from the innermost block of $T(7a+1,7)$.

Either of the previous lemmas imply the following lower bound, which is also implicit in the work of Hayashi *et al.* [37].

Theorem 2.4. *For every positive integer* n*, there are closed curves with* n *vertices whose defects are* $n^{3/2}/3 - O(n)$ *and* $-n^{3/2}/3 + O(n)$*, and therefore requires at least* $n^{3/2}/6 - O(n)$ *homotopy moves to reduce to a simple closed curve.*

Proof: The lower bound follows from the previous lemmas by setting $a = 1$. If n is a prefect square, then the flat torus knot $T(\sqrt{n}+1,\sqrt{n})$ has n vertices and defect $-2\binom{\sqrt{n}}{3}$. If n is not a perfect square, we can achieve defect $-2\binom{\lfloor\sqrt{n}\rfloor}{3}$ by applying 0→1 moves to the curve $T(\lfloor\sqrt{n}\rfloor+1,\lfloor\sqrt{n}\rfloor)$. Similarly, we obtain an n-vertex curve with defect $2\binom{\lfloor\sqrt{n+1}\rfloor+1}{3}$ by adding loops to the curve $T(\lfloor\sqrt{n+1}\rfloor,\lfloor\sqrt{n+1}\rfloor+1)$. Lemma 2.1 now immediately implies the lower bound on homotopy moves. □

2.3 Multicurves

Our previous results immediately imply that simplifying a multicurve with n vertices requires at least $\Omega(n^{3/2})$ homotopy moves; in this section we derive additional lower bounds in terms of the number

of constituent curves. We distinguish between two natural variants of simplification: transforming a multicurve into an *arbitrary* set of disjoint simple closed curves, or into a *particular* set of disjoint simple closed curves.

Both lower bound proofs rely on the classical notion of ***winding number***. Let γ be an arbitrary closed curve in the plane, let p be any point outside the image of γ, and let ρ be any ray from p to infinity that intersects γ transversely. The winding number of γ around p, which we denote $\boldsymbol{wind(\gamma, p)}$, is the number of times γ crosses ρ from right to left, minus the number of times γ crosses ρ from left to right. The winding number does not depend on the particular choice of ray ρ. All points in the same face of γ have the same winding number. Moreover, if there is a homotopy from one curve γ to another curve γ', where the image of any intermediate curve does not include p, then $wind(\gamma, p) = wind(\gamma', p)$ [38].

Lemma 2.5. *Transforming a k-curve with n vertices in the plane into k arbitrary disjoint circles requires $\Omega(nk)$ homotopy moves in the worst case.*

Proof: For arbitrary positive integers n and k, we construct a multicurve with k disjoint constituent curves, all but one of which are simple, as follows. The first $k-1$ constituent curves $\gamma_1, \ldots, \gamma_{k-1}$ are disjoint circles inside the open unit disk centered at the origin. (The precise configuration of these circles is unimportant.) The remaining constituent curve γ_o is a spiral winding $n+1$ times around the closed unit disk centered at the origin, plus a line segment connecting the endpoints of the spiral; γ_o is the simplest possible curve with winding number $n+1$ around the origin. Let γ be the disjoint union of these k curves; we claim that $\Omega(nk)$ homotopy moves are required to simplify γ. See Figure 2.6.

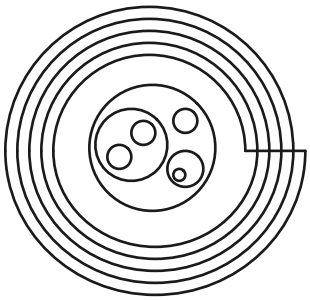

Figure 2.6. Simplifying this multicurve requires $\Omega(nk)$ homotopy moves.

Consider the faces of the outer curve γ_o during any homotopy of γ. Adjacent faces of γ_o have winding numbers that differ by 1, and the outer face has winding number 0. Thus, for any non-negative integer w, as long as the maximum absolute winding number $\left|\max_p wind(\gamma_o, p)\right|$ is at least w, the curve γ_o has at least $w+1$ faces (including the outer face) and therefore at least $w-1$ vertices, by Euler's formula. On the other hand, if any curve γ_i intersects a face of γ_o, no homotopy move can remove that face until the intersection between γ_i and γ_o is removed. Thus, before the simplification of γ_o is complete, each curve γ_i must intersect only faces with winding number 0, 1, or -1.

For each index i, let w_i denote the maximum absolute winding number of γ_o around any point of γ_i:

$$w_i := \max_\theta \left| wind\left(\gamma_o, \gamma_i(\theta)\right)\right|.$$

Let $W = \sum_i w_i$. Initially, $W = k(n+1)$, and when γ_o first becomes simple, we must have $W \le k$. Each homotopy move changes W by at most 1; specifically, at most one term w_i changes at all, and that term either increases or decreases by 1. The $\Omega(nk)$ lower bound now follows immediately. □

Theorem 2.6. *Transforming a k-curve with n vertices in the plane into an arbitrary set of k simple closed curves requires $\Omega(n^{3/2} + nk)$ homotopy moves in the worst case.*

We say that a collection of k disjoint simple closed curves is ***nested*** if some point lies in the interior of every curve, and ***unnested*** if the curves have disjoint interiors.

Lemma 2.7. *Transforming k nested circles in the plane into k unnested circles requires $\Omega(k^2)$ homotopy moves.*

Proof: Let γ and γ' be two nested circles, with γ' in the interior of γ and with γ directed counterclockwise. Suppose we apply an arbitrary homotopy to these two curves. If the curves remain disjoint during the entire homotopy, then γ' always lies inside a face of γ with winding number 1; in short, the two curves remain nested. Thus, any sequence of homotopy moves that takes γ and γ' to two non-nested simple closed curves contains at least one 0→2 move that makes the curves cross (and symmetrically at least one 2→0 move that makes them disjoint again).

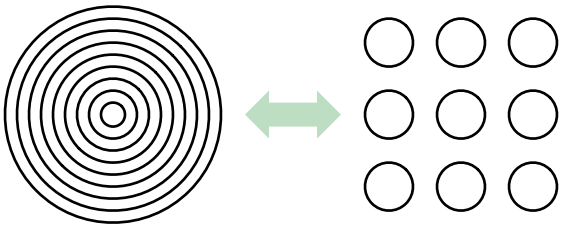

Figure 2.7. Nesting or unnesting k circles requires $\Omega(k^2)$ homotopy moves.

Consider a set of k nested circles. Each of the $\binom{k}{2}$ pairs of circles requires at least one 0→2 move and one 2→0 move to unnest. Because these moves involve distinct pairs of curves, at least $\binom{k}{2}$ 0→2 moves and $\binom{k}{2}$ 2→0 moves, and thus at least $k^2 - k$ moves altogether, are required to unnest every pair. □

Theorem 2.8. *Transforming a k-curve with n vertices in the plane into k nested (or unnested) circles requires $\Omega(n^{3/2} + nk + k^2)$ homotopy moves in the worst case.*

Corollary 2.9. *Transforming one k-curve with at most n vertices into another k-curve with at most n vertices requires $\Omega(n^{3/2} + nk + k^2)$ homotopy moves in the worst case.*

Although our lower bound examples consist of disjoint curves, all of these lower bounds apply without modification to *connected* multicurves, because any k-curve can be connected with at most $k - 1$ 0→2 moves. On the other hand, any connected k-curve has at least $2k - 2$ vertices, so the $\Omega(k^2)$ terms in Theorem 2.8 and Corollary 2.9 are redundant.

3 Electrical Transformations

Now we consider a related set of local operations on plane graphs, called ***facial electrical transformations***, consisting of six operations in three dual pairs, as shown in Figure 3.1.

- *degree-1 reduction*: Contract the edge incident to a vertex of degree 1, or delete the edge incident to a face of degree 1
- *series-parallel reduction*: Contract either edge incident to a vertex of degree 2, or delete either edge incident to a face of degree 2
- *ΔY transformation*: Delete a vertex of degree 3 and connect its neighbors with three new edges, or delete the edges bounding a face of degree 3 and join the vertices of that face to a new vertex.

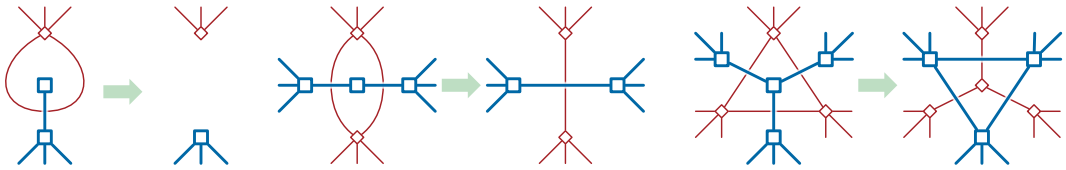

Figure 3.1. Facial electrical transformations in a plane graph G and its dual graph G^*.

Electrical transformations are usually defined more generally as a set of operations performed on abstract graphs, which have been used since the end of the 19th century [40, 53] to analyze resistor networks and other electrical circuits, but have since been applied to a number of other combinatorial problems on planar graphs, including shortest paths and maximum flows [2]; multicommodity flows [23]; and counting spanning trees, perfect matchings, and cuts [16]. We refer to our earlier preprint [11, Section 1.1] for a more detailed history and an expanded list of applications. However, all the algorithms we describe below reduce any plane graph to a single vertex using only *facial* electrical transformations as defined above.

In light of these applications, it is natural to ask *how many* facial electrical transformations are required in the worst case. The earliest algorithm for reducing a plane graph to a single vertex already follows from Steinitz's bigon-reduction argument, which we described in the introduction [55, 56]. Steinitz reduced local transformations of plane *graphs* to local transformations of planar *curves* by defining the *medial graphs* ("Θ-Prozess"), which we consider in detail below. Later algorithms were given by Truemper [58], Feo and Provan [24], and others. Both Steinitz's algorithm and Feo and Provan's algorithm require at most $O(n^2)$ facial electrical transformations; this is the best upper bound known.

Even the special case of regular grids is open and interesting. Truemper [58, 59] describes a method to reduce the $p \times p$ grid in $O(p^3)$ steps. Nakahara and Takahashi [45] prove an upper bound of $O(\min\{pq^2, p^2q\})$ for the $p \times q$ cylindrical grid. Because every n-vertex plane graph is a minor of an $O(n) \times O(n)$ grid [57, 60], both of these results imply an $O(n^3)$ upper bound for arbitrary plane graphs; see Lemma 3.1. Feo and Provan [24] claim without proof that Truemper's algorithm actually performs only $O(n^2)$ electrical transformations. On the other hand, the smallest (cylindrical) grid containing every n-vertex plane graph as a minor has size $\Omega(n) \times \Omega(n)$ [60]. Archdeacon *et al.* [7] asked whether the $O(n^{3/2})$ upper bound for square grids can be improved to near-linear:

> It is possible that a careful implementation and analysis of the grid-embedding schemes can lead to an $O(n\sqrt{n})$-time algorithm for the general planar case. It would be interesting to obtain a near-linear algorithm for the grid.... However, it may well be that reducing planar grids is $\Omega(n\sqrt{n})$.

Most of these earlier algorithms actually solve a more difficult problem, first considered by Akers [2] and Lehman [43] and later solved by Epifanov [21], of reducing a planar graph with two special vertices called *terminals* to a single edge between the two terminals. In this context, any electrical transformation that contracts an edge incident to a terminal is forbidden. Unfortunately, not every two-terminal plane graph can be reduced to a single edge using only facial electrical transformations; Figure 3.2 shows two examples. However, it is sufficient to allow loop reductions, parallel reductions, and $\Delta \to Y$ transformations to be performed on faces that contain a terminal vertex of degree 1 (and nothing else) [24]. It is an open question whether our lower bound still holds if these additional non-facial transformations are allowed.

3.1 Definitions

The ***medial graph*** of a plane graph G, which we denote $\boldsymbol{G^\times}$, is another plane graph whose vertices correspond to the edges of G and whose edges correspond to incidences (with multiplicity) between

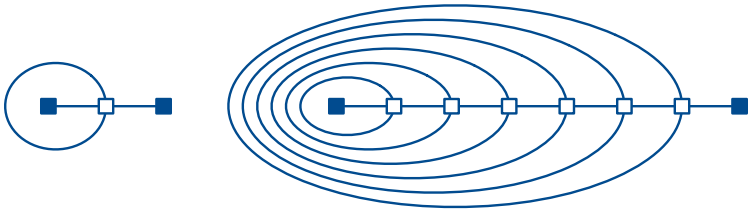

Figure 3.2. Plane graphs with two terminals that cannot be further reduced using only facial electrical transformations.

vertices of G and faces of G. Two vertices of $G^\times$ are connected by an edge if and only if the corresponding edges in G are consecutive in cyclic order around some vertex, or equivalently, around some face in G. Every vertex in every medial graph has degree 4; thus, every medial graph is the image of a multicurve. The medial graphs of any plane graph G and its dual G^* are identical. To avoid trivial boundary cases, we define the medial graph of an isolated vertex to be a circle.

Facial electrical transformations in any plane graph G correspond to local transformations in the medial graph $G^\times$ that are almost identical to homotopy moves. Each degree-1 reduction in G corresponds to a $1\to0$ homotopy move in $G^\times$, and each ΔY transformation in G corresponds to a $3\to3$ homotopy move in $G^\times$. A series-parallel reduction in G contracts an empty bigon in $G^\times$ to a single vertex. Extending our earlier notation, we call this transformation a **2→1** move. We collectively refer to these transformations and their inverses as ***medial electrical moves***; see Figure 3.3.

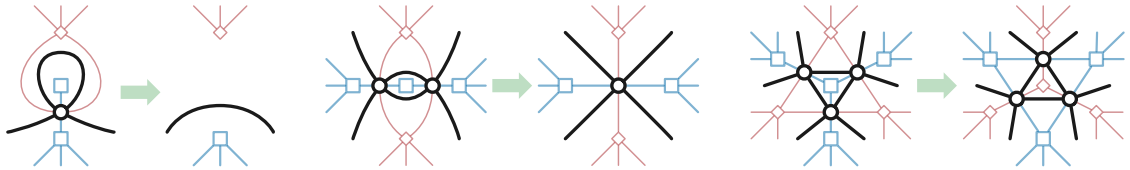

Figure 3.3. Medial electrical moves $1\to0$, $2\to1$, and $3\to3$.

Smoothing a multicurve γ at a vertex x means replacing the intersection of γ with a small neighborhood of x with two disjoint simple paths, so that the result is another multicurve. (There are two possible smoothings at each vertex; see Figure 3.4.) A ***smoothing*** of γ is any graph obtained by smoothing zero or more vertices of γ, and a ***proper smoothing*** of γ is any smoothing other than γ itself. For any plane graph G, the (proper) smoothings of the medial graph $G^\times$ are precisely the medial graphs of (proper) minors of G.

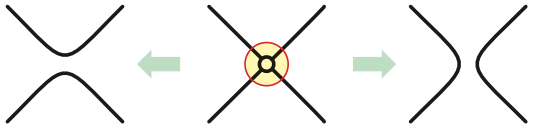

Figure 3.4. Smoothing a vertex.

3.2 Electrical to Homotopy

The main result of this section is that the number of *homotopy* moves required to simplify a closed curve is a lower bound on the number of *medial electrical moves* required to simplify the same closed curve. This result is already implicit in the work of Noble and Welsh [46], and most of our proofs closely follow theirs. We include the proofs here to make the inequalities explicit and to keep the paper self-contained.

For any connected multicurve (or 4-regular plane graph) γ, let $\boldsymbol{X(\gamma)}$ denote the minimum number of medial electrical moves required to reduce γ to a simple closed curve, and let $\boldsymbol{H(\gamma)}$ is the minimum number of homotopy moves required to reduce γ to an arbitrary collection of disjoint simple closed curves.

The following key lemma follows from close reading of proofs by Truemper [58, Lemma 4] and several others [7, 27, 45, 46] that every minor of a ΔY-reducible graph is also ΔY-reducible. Our proof most closely resembles an argument of Gitler [27, Lemma 2.3.3], but restated in terms of medial electrical moves to simplify the case analysis.

Lemma 3.1. *For any connected plane graph G, reducing any connected proper minor of G to a single vertex requires strictly fewer facial electrical transformations than reducing G to a single vertex. Equivalently, $X(\overline{\gamma}) < X(\gamma)$ for every connected proper smoothing $\overline{\gamma}$ of every connected multicurve γ.*

Proof: Let γ be a connected multicurve, and let $\overline{\gamma}$ be a connected proper smoothing of γ. If γ is already simple, the lemma is vacuously true. Otherwise, the proof proceeds by induction on $X(\gamma)$.

We first consider the special case where $\overline{\gamma}$ is obtained from γ by smoothing a single vertex x. Let γ' be the result of the first medial electrical move in the minimum-length sequence that reduces γ to a simple closed curve. We immediately have $X(\gamma) = X(\gamma') + 1$. There are two nontrivial cases to consider.

First, suppose the move from γ to γ' does not involve the smoothed vertex x. Then we can apply the same move to $\overline{\gamma}$ to obtain a new graph $\overline{\gamma}'$; the same graph can also be obtained from γ' by smoothing x. We immediately have $X(\overline{\gamma}) \le X(\overline{\gamma}')+1$, and the inductive hypothesis implies $X(\overline{\gamma}')+1 < X(\gamma')+1 = X(\gamma)$.

Now suppose the first move in Σ does involve x. In this case, we can apply at most one medial electrical move to $\overline{\gamma}$ to obtain a (possibly trivial) smoothing $\overline{\gamma}'$ of γ'. There are eight subcases to consider, shown in Figure 3.5. One subcase for the 0→1 move is impossible, because $\overline{\gamma}$ is connected. In the remaining 0→1 subcase and one 2→1 subcase, the curves $\overline{\gamma}$, $\overline{\gamma}'$ and γ' are all isomorphic, which implies $X(\overline{\gamma}) = X(\overline{\gamma}') = X(\gamma') = X(\gamma) - 1$. In all remaining subcases, $\overline{\gamma}'$ is a connected proper smoothing of γ', so the inductive hypothesis implies $X(\overline{\gamma}) \le X(\overline{\gamma}') + 1 < X(\gamma') + 1 = X(\gamma)$.

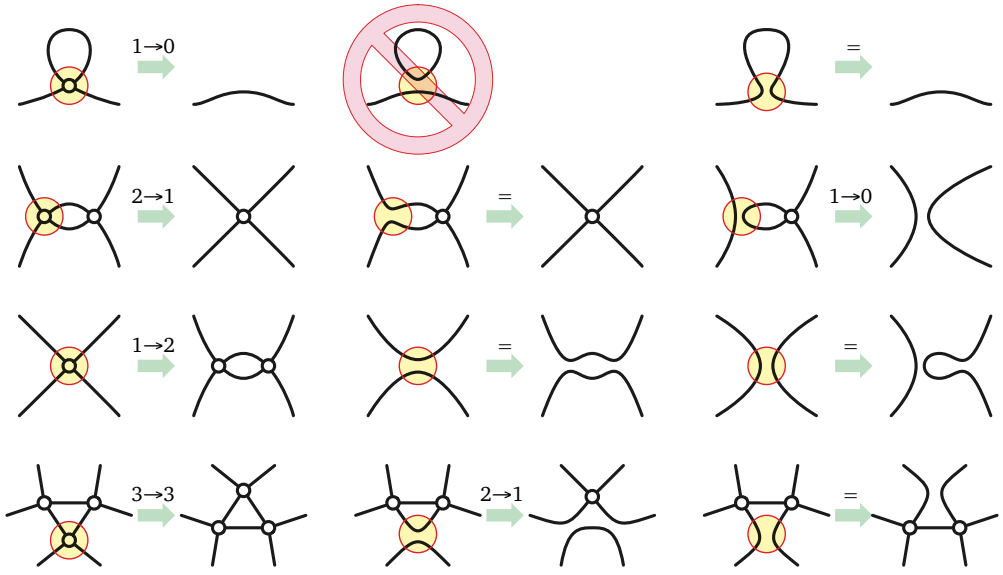

Figure 3.5. Cases for the proof of the Lemma 3.1; the circled vertex is x.

Finally, we consider the more general case where $\overline{\gamma}$ is obtained from γ by smoothing more than one vertex. Let $\widetilde{\gamma}$ be any intermediate curve, obtained from γ by smoothing just one of the vertices that were

smoothed to obtain $\overline{\gamma}$. As $\overline{\gamma}$ is a connected smoothing of $\widetilde{\gamma}$, the curve $\widetilde{\gamma}$ itself must be connected too. Our earlier argument implies that $X(\widetilde{\gamma}) < X(\gamma)$. Thus, the inductive hypothesis implies $X(\overline{\gamma}) < X(\widetilde{\gamma})$, which completes the proof. □

Lemma 3.2. *For every connected multicurve γ, there is a minimum-length sequence of medial electrical moves that reduces γ to a simple closed curve and that does not contain* 0→1 *or* 1→2 *moves.*

Proof: Our proof follows an argument of Noble and Welsh [46, Lemma 3.2].

Consider a minimum-length sequence of medial electrical moves that reduces an arbitrary connected multicurve γ to a simple closed curve. For any integer $i \geq 0$, let γ_i denote the result of the first i moves in this sequence; in particular, $\gamma_0 = \gamma$ and $\gamma_{X(\gamma)}$ is a simple closed curve. Minimality of the reduction sequence implies that $X(\gamma_i) = X(\gamma) - i$ for all i. Now let i be an arbitrary index such that γ_i has one more vertex than γ_{i-1}. Then γ_{i-1} is a connected proper smoothing of γ_i, so Lemma 3.1 implies that $X(\gamma_{i-1}) < X(\gamma_i)$, giving us a contradiction. □

Lemma 3.3. $X(\gamma) \geq H(\gamma)$ *for every closed curve γ.*

Proof: The proof proceeds by induction on $X(\gamma)$, following an argument of Noble and Welsh [46, Proposition 3.3].

Let γ be a closed curve. If $X(\gamma) = 0$, then γ is already simple, so $H(\gamma) = 0$. Otherwise, let Σ be a minimum-length sequence of medial electrical moves that reduces γ to a simple closed curve. Lemma 3.2 implies that we can assume that the first move in Σ is neither 0→1 nor 1→2. If the first move is 1→0 or 3→3, the theorem immediately follows by induction.

The only interesting first move is 2→1. Let γ' be the result of this 2→1 move, and let $\overline{\gamma}$ be the result of the corresponding 2→0 homotopy move. The minimality of Σ implies that $X(\gamma) = X(\gamma') + 1$, and we trivially have $H(\gamma) \leq H(\overline{\gamma}) + 1$. Because γ consists of *one* single curve, $\overline{\gamma}$ is also a single curve and is therefore connected. The curve $\overline{\gamma}$ is also a proper smoothing of γ', so the Lemma 3.1 implies $X(\overline{\gamma}) < X(\gamma') < X(\gamma)$. Finally, the inductive hypothesis implies that $X(\overline{\gamma}) \geq H(\overline{\gamma})$, which completes the proof. □

We call a plane graph G ***unicursal*** if its medial graph $G^\times$ is the image of a single closed curve.

Theorem 3.4. *For every connected plane graph G and every unicursal minor H of G, reducing G to a single vertex requires at least $|defect(H^\times)|/2$ facial electrical transformations.*

Proof: Either H equals G, or Lemma 3.1 implies that reducing a proper minor H of G to a single vertex requires strictly fewer facial electrical transformations than reducing G to a single vertex. Note that facial electrical transformations performed on H corresponds precisely to medial electrical moves performed on $H^\times$. Now because $\gamma := H^\times$ is unicursal, Lemma 2.1 and Lemma 3.3 implies that $X(\gamma) \geq H(\gamma) \geq |defect(\gamma)|/2$. □

3.3 Cylindrical Grids

Finally, we derive explicit lower bounds for the number of facial electrical transformations required to reduce any cylindrical grid to a single vertex. For any positive integers p and q, we define two cylindrical grid graphs; see Figure 3.6.

- $\boldsymbol{C(p,q)}$ is the Cartesian product of a cycle of length q and a path of length $p-1$. If q is odd, then the medial graph of $C(p,q)$ is the flat torus knot $T(2p,q)$.

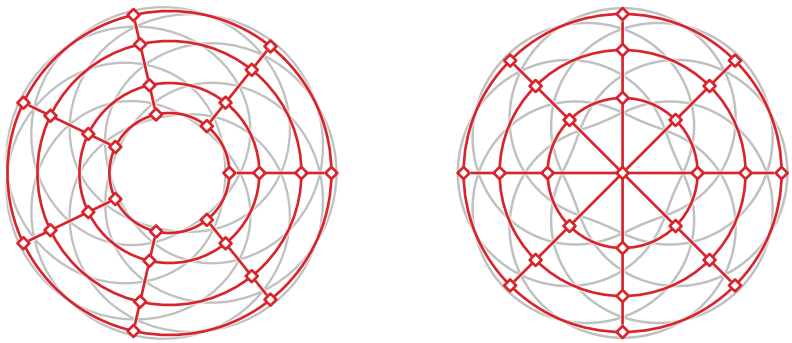

Figure 3.6. The cylindrical grid graphs $C(4,7)$ and $C'(3,8)$ and (in light gray) their medial graphs $T(8,7)$ and $T(7,8)$.

- $\boldsymbol{C'(p,q)}$ is obtained by connecting a new vertex to the vertices of one of the q-gonal faces of $C(p,q)$, or equivalently, by contracting one of the q-gonal faces of $C(p+1,q)$ to a single vertex. If q is even, then the medial graph of $C'(p,q)$ is the flat torus knot $T(2p+1,q)$.

Corollary 3.5. *For all positive integers p and q, the cylindrical grid $C(p,q)$ requires $\Omega(\min\{p^2q, pq^2\})$ facial electrical transformations to reduce to a single vertex.*

Proof: First suppose $p \le q$. Because $C(p-1,q)$ is a minor of $C(p,q)$, we can assume without loss of generality that p is even and $p < q$. Let H denote the cylindrical grid $C(p/2, ap+1)$, where $a := \lfloor (q-1)/p \rfloor \ge 1$. H is a minor of $C(p,q)$ (because $ap+1 \le q$), and the medial graph of H is the flat torus knot $T(p, ap+1)$. Lemma 2.2 implies

$$\mathit{defect}\Big(T(p,ap+1)\Big) = 2a\binom{p+1}{3} = \Omega(ap^3) = \Omega(p^2q).$$

Theorem 3.4 now implies that reducing $C(p,q)$ requires at least $\Omega(p^2q)$ facial electrical transformations.

The symmetric case $p > q$ is similar. We can assume without loss of generality that q is odd. Let H denote the cylindrical grid $C'(aq,q)$, where $a := \lfloor (p-1)/q \rfloor \ge 1$. H is a proper minor of $C(p,q)$ (because $aq < p$), and the medial graph of H is the flat torus knot $T(2aq+1,q)$. Lemma 2.3 implies

$$\Big|\mathit{defect}\Big(T(2aq+1,q)\Big)\Big| = 4a\binom{q}{3} = \Omega(aq^3) = \Omega(pq^2).$$

Theorem 3.4 now implies that reducing $C(p,q)$ requires at least $\Omega(pq^2)$ facial electrical transformations. □

In particular, reducing any $\Theta(\sqrt{n}) \times \Theta(\sqrt{n})$ cylindrical grid requires at least $\Omega(n^{3/2})$ facial electrical transformations. Our lower bound matches an $O(\min\{pq^2, p^2q\})$ upper bound by Nakahara and Takahashi [45]. Because every $p \times q$ rectangular grid contains $C(\lfloor p/3 \rfloor, \lfloor q/3 \rfloor)$ as a minor, the same $\Omega(\min\{p^2q, pq^2\})$ lower bound applies to rectangular grids. In particular, Truemper's $O(p^3) = O(n^{3/2})$ upper bound for the $p \times p$ square grid [58] is tight. Finally, because every plane graph with treewidth t contains an $\Omega(t) \times \Omega(t)$ grid minor [52], reducing *any* n-vertex plane graph with treewidth t requires at least $\Omega(t^3 + n)$ facial electrical transformations.

Like Gitler [27], Feo and Provan [24], and Archdeacon *et al.* [7], we conjecture that any n-vertex plane graph can be reduced to a vertex using $O(n^{3/2})$ facial electrical transformations. More ambitiously, we conjecture an upper bound of $O(nt)$ for any n-vertex plane graph with treewidth t.

4 Upper Bound

For any point p, let ***depth*(p,γ)** denote the minimum number of times a path from p to infinity crosses γ. Any two points in the same face of γ have the same depth, so each face f has a well-defined depth, which is its distance to the outer face in the dual graph of γ; see Figure 4.1. The depth of the curve, denoted ***depth*(γ)**, is the maximum depth of the faces of γ; and the ***potential* $D(\gamma)$** is the sum of the depths of the faces. Euler's formula implies that any 4-regular plane graph with n vertices has exactly $n+2$ faces; thus, for any curve γ with n vertices, we have $n+1 \le D(\gamma) \le (n+1)\cdot \mathit{depth}(\gamma)$.

4.1 Contracting Simple Loops

Lemma 4.1. *Every closed curve γ in the plane can be simplified using at most $3D(\gamma)-3$ homotopy moves.*

Proof: We prove the statement by induction on the number of vertices in γ. The lemma is trivial if γ is already simple, so assume otherwise. Let $x := \gamma(\theta) = \gamma(\theta')$ be the first vertex to be visited twice by γ after the (arbitrarily chosen) basepoint $\gamma(0)$. Let α denote the subcurve of γ from $\gamma(\theta)$ to $\gamma(\theta')$; our choice of x implies that α is a simple loop. Let m and s denote the number of vertices and maximal subpaths of γ in the interior of α respectively.

Finally, let γ' denote the closed curve obtained from γ by removing α. The first stage of our algorithm transforms γ into γ' by contracting the loop α via homotopy moves.

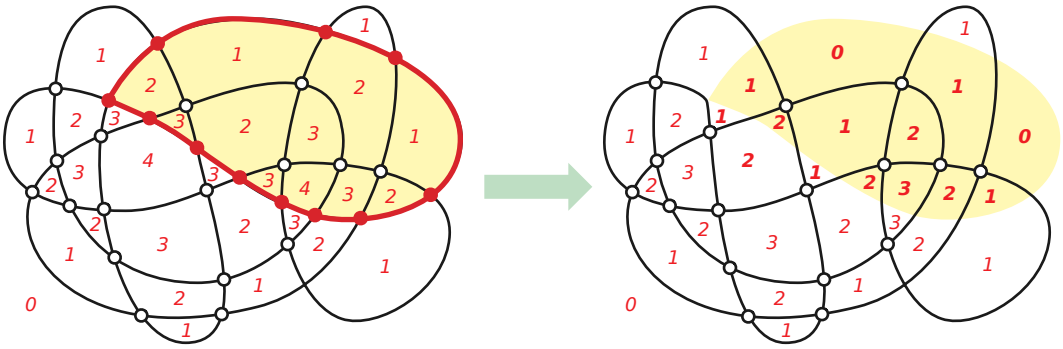

Figure 4.1. Transforming γ into γ' by contracting a simple loop. Numbers are face depths.

We remove the vertices and edges from the interior of α one at a time as follows (see Figure 4.2). If we can perform a 2→0 move to remove one edge of γ from the interior of α and decrease s, we do so. Otherwise, either α is empty, or some vertex of γ lies inside α. In the latter case, at least one vertex x inside α has a neighbor that lies on α. We move x outside α with a 0→2 move (which increases s by 1) followed by a 3→3 move (which decreases m by 1). Once α is an empty loop, we remove it with a single 1→0 move. Altogether, our algorithm transforms γ into γ' using at most $3m+s+1$ homotopy moves. Let M denote the actual number of homotopy moves used.

Euler's formula implies that α contains exactly $m+s+1$ faces of γ. The Jordan curve theorem implies that $\mathit{depth}(p,\gamma') \le \mathit{depth}(p,\gamma)-1$ for any point p inside α, and trivially $\mathit{depth}(p,\gamma') \le \mathit{depth}(p,\gamma)$ for any point p outside α. It follows that $D(\gamma') \le D(\gamma) - (m+s+1) \le D(\gamma) - M/3$, and therefore $M \le 3D(\gamma) - 3D(\gamma')$. The induction hypothesis implies that we can recursively simplify γ' using at most $3D(\gamma')-3$ moves. The lemma now follows immediately. □

Our upper bound is a factor of 3 larger than Feo and Provan's [24]; however our algorithm has the advantage that it extends to *tangles*, as described in the next subsection.

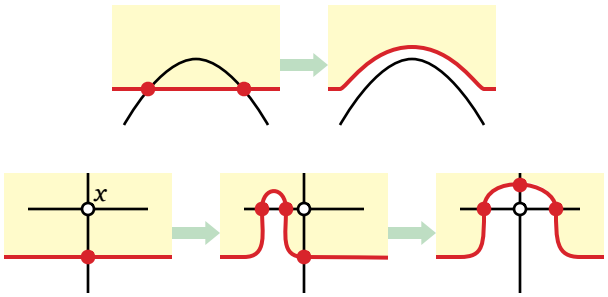

Figure 4.2. Moving a loop over an interior empty bigon or an interior vertex.

4.2 Tangles

A ***tangle*** is a collection of boundary-to-boundary paths $\gamma_1, \gamma_2, \ldots, \gamma_s$ in a closed topological disk Σ, which (self-)intersect only pairwise, transversely, and away from the boundary of Σ. This terminology is borrowed from knot theory, where a tangle usually refers to the intersection of a knot or link with a closed 3-dimensional ball [14, 18]; our tangles are perhaps more properly called *flat tangles*, as they are images of tangles under appropriate projection. (Our tangles are unrelated to the obstructions to small branchwidth introduced by Robertson and Seymour [51].) Transforming a curve into a tangle is identical to (an inversion of) the *flarb* operation defined by Allen *et al.* [5].

We call each individual path γ_i a ***strand*** of the tangle. The ***boundary*** of a tangle is the boundary of the disk Σ that contains it; we usually denote the boundary by σ. By the Jordan-Schönflies theorem, we can assume without loss of generality that σ is actually a circle. We can obtain a tangle from any closed curve γ by considering its restriction to any closed disk whose boundary σ intersects γ transversely away from its vertices; we call this restriction the ***interior tangle*** of σ.

The strands and boundary of any tangle define a plane graph T whose boundary vertices each have degree 3 and whose interior vertices each have degree 4. Depths and potential of a tangle are defined exactly as for closed curves: The depth of any face f of T is its distance to the outer face in the dual graph T^*; the depth of the tangle is its maximum face depth; and the potential $D(T)$ of the tangle is the sum of its face depths.

A tangle is ***tight*** if every pair of strands intersects at most once and ***loose*** otherwise. Every loose tangle contains either an empty loop or a (not necessarily empty) bigon. Thus, any tangle with n vertices can be transformed into a tight tangle—or less formally, *tightened*—in $O(n^2)$ homotopy moves using Steinitz's algorithm. On the other hand, there are infinite classes of loose tangles for which no homotopy move that decreases the potential, so we cannot directly apply Feo and Provan's algorithm to this setting.

We describe a two-phase algorithm to tighten any tangle. First, we remove any self-intersections in the individual strands, by contracting loops as in the proof of Lemma 4.1. Once each strand is simple, we move the strands so that each pair intersects at most once. See Figure 4.3.

Lemma 4.2. *Every tangle T with n vertices and s strands, such that every strand is simple, can be tightened using at most $3ns$ homotopy moves.*

Proof: We prove the lemma by induction on s. The base case when $s = 1$ is trivial, so assume $s \geq 2$.

Fix an arbitrary reference point on the boundary circle σ that is not an endpoint of a strand. For each index i, let σ_i be the arc of σ between the endpoints of γ_i that does not contain the reference point. A strand γ_i is *extremal* if the corresponding arc σ_i does not contain any other arc σ_j.

Choose an arbitrary extremal strand γ_i. Let m_i denote the number of tangle vertices in the interior of the disk bounded by γ_i and the boundary arc σ_i; call this disk Σ_i. Let s_i denote the number of intersections between γ_i and other strands. Finally, let γ_i' be a path inside the disk Σ defining tangle T,

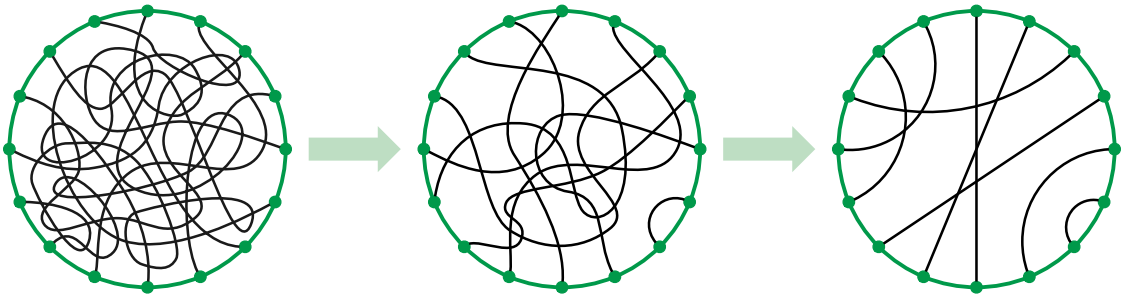

Figure 4.3. Tightening a tangle in two phases: First simplifying the individual strands, then removing excess crossings between pairs of strands.

with the same endpoints as γ_i, that intersects each other strand in T at most once, such that the disk bounded by σ_i and γ_i' has no tangle vertices inside its interior. (See Figure 4.4 for an illustration; the red strand in the left tangle is γ_i, the red strand in the middle tangle is γ_i', and the shaded disk is Σ_i.).

We can deform γ_i into γ_i' using essentially the algorithm from Lemma 4.1; the disk Σ_i is contracted along with γ_i in the process. If Σ_i contains an empty bigon with one side in γ_i, remove it with a $2\to0$ move (which decreases s_i by 1). If Σ_i has an interior vertex with a neighbor on γ_i, remove it using at most two homotopy moves (which increases s_i by 1 and decreases m_i by 1). Altogether, this deformation requires at most $3m_i + s_i \le 3n$ homotopy moves.

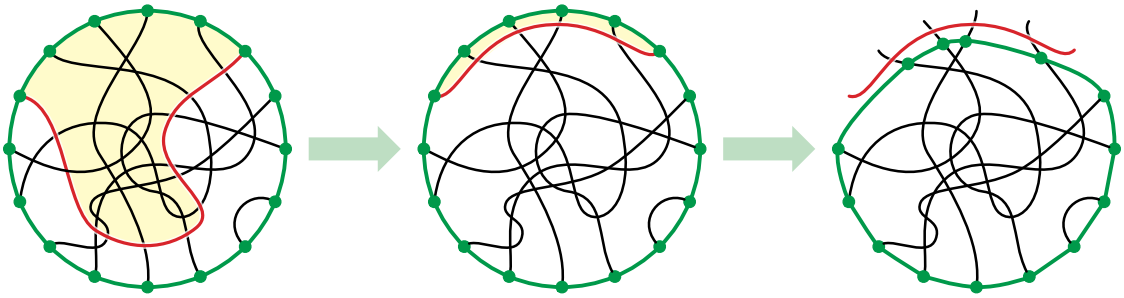

Figure 4.4. Moving one strand out of the way and shrinking the tangle boundary.

After deforming γ_i to γ_i', we redefine the tangle by "shrinking" its boundary curve slightly to exclude γ_i', without creating or removing any vertices in the tangle or endpoints on the boundary (see the right of Figure 4.4). We emphasize that shrinking the boundary does not modify the strands and therefore does not require any homotopy moves. The resulting smaller tangle has exactly $s-1$ strands, each of which is simple. Thus, the induction hypothesis implies that we can recursively tighten this smaller tangle using at most $3n(s-1)$ homotopy moves. □

Corollary 4.3. *Every tangle T with n vertices and s strands can be tightened using at most $3D(T)+3ns$ homotopy moves.*

Proof: As long as T contains at least one non-simple strand, we identify a simple loop α in that strand and remove it as described in the proof of Lemma 4.1. Suppose there are m vertices and t maximal subpaths in the interior of α, and let M be the number of homotopy moves required to remove α. The algorithm in the proof of Lemma 4.1 implies that $M \le 3m + t + 1$, and Euler's formula implies that α contains $m + t + 1 \ge M/3$ faces. Removing α decreases the depth of each of these faces by at least 1 and therefore decreases the potential of the tangle by at least $M/3$.

Let T' be the remaining tangle after all such loops are removed. Our potential analysis for a single loop implies inductively that transforming T into T' requires at most $3D(T)-3D(T') \le 3D(T)$ homotopy

moves. Because T' still has s strands and at most n vertices, Lemma 4.2 implies that we can tighten T' with at most $3ns$ additional homotopy moves. □

4.3 Main Algorithm

We call a simple closed curve σ ***useful*** for γ if σ intersects γ transversely away from its vertices, and the interior tangle T of σ has at least s^2 vertices, where $s := |\sigma \cap \gamma|/2$ is the number of strands. Our main algorithm repeatedly finds a useful closed curve whose interior tangle has $O(\sqrt{n})$ depth, and tightens its interior tangle; if there are no useful closed curves, then we fall back to the loop-contraction algorithm of Lemma 4.1.

Lemma 4.4. *Let γ be an arbitrary non-simple closed curve in the plane with n vertices. Either there is a useful simple closed curve for γ whose interior tangle has depth $O(\sqrt{n})$, or the depth of γ is $O(\sqrt{n})$.*

Proof: To simplify notation, let $d := \mathit{depth}(\gamma)$. For each integer j between 1 and d, let R_j be the set of points p with $\mathit{depth}(p,\gamma) \geq d+1-j$, and let $\tilde{R}_j$ denote a small open neighborhood of the closure of $R_j \cup \tilde{R}_{j-1}$, where $\tilde{R}_0$ is the empty set. Each region $\tilde{R}_j$ is the disjoint union of closed disks, whose boundary cycles intersect γ transversely away from its vertices, if at all. In particular, $\tilde{R}_d$ is a disk containing the entire curve γ.

Fix a point z such that $\mathit{depth}(z,\gamma) = d$. For each integer j, let Σ_j be the unique component of $\tilde{R}_j$ that contains z, and let σ_j be the boundary of Σ_j. Then $\sigma_1, \sigma_2, \ldots, \sigma_d$ are disjoint, nested, simple closed curves; see Figure 4.5. Let n_j be the number of vertices and let $s_j := |\gamma \cap \sigma_j|/2$ be the number of strands of the interior tangle of σ_j. For notational convenience, we define $\Sigma_0 := \emptyset$ and thus $n_0 = s_0 = 0$. We ignore the outermost curve σ_d, because it contains the entire curve γ. The next outermost curve σ_{d-1} contains every vertex of γ, so $n_{d-1} = n$.

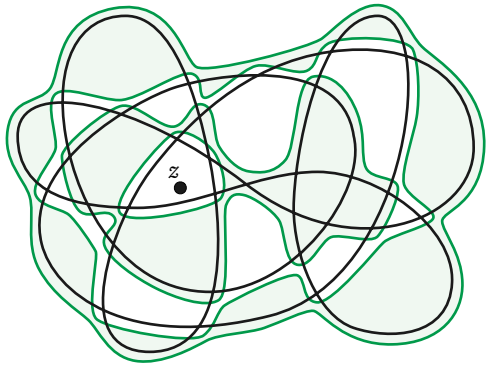

Figure 4.5. Nested depth cycles around a point of maximum depth.

By construction, for each j, the interior tangle of σ_j has depth $j+1$. Thus, to prove the lemma, it suffices to show that if none of the curves $\sigma_1, \sigma_2, \ldots, \sigma_{d-1}$ is useful, then $d = O(\sqrt{n})$.

Fix an index j. Each edge of γ crosses σ_j at most twice. Any edge of γ that crosses σ_j has at least one endpoint in the annulus $\Sigma_j \setminus \Sigma_{j-1}$, and any edge that crosses σ_j twice has both endpoints in $\Sigma_j \setminus \Sigma_{j-1}$. Conversely, each vertex in Σ_j is incident to at most two edges that cross σ_j and no edges that cross σ_{j+1}. It follows that $|\sigma_j \cap \gamma| \leq 2(n_j - n_{j-1})$, and therefore $n_j \geq n_{j-1} + s_j$. Thus, by induction, we have

$$n_j \geq \sum_{i=1}^{j} s_i$$

for every index j.

Now suppose no curve σ_j with $1 \le j < d$ is useful. Then we must have $s_j^2 > n_j$ and therefore

$$s_j^2 > \sum_{i=1}^{j} s_i$$

for all $1 \le j < d$. Trivially, $s_1 \ge 1$, because γ is non-simple. A straightforward induction argument implies that $s_j \ge (j+1)/2$ and therefore

$$n \;=\; n_{d-1} \;\ge\; \sum_{i=1}^{d-1} \frac{i+1}{2} \;\ge\; \frac{1}{2}\binom{d+1}{2} \;>\; \frac{d^2}{4}.$$

We conclude that $d \le 2\sqrt{n}$, which completes the proof. □

Theorem 4.5. *Every closed curve in the plane with n vertices can be simplified in $O(n^{3/2})$ homotopy moves.*

Proof: Let γ be an arbitrary closed curve in the plane with n vertices. If γ has depth $O(\sqrt{n})$, Lemma 4.1 and the trivial upper bound $D(\gamma) \le (n+1) \cdot \mathit{depth}(\gamma)$ imply that we can simplify γ in $O(n^{3/2})$ homotopy moves. For purposes of analysis, we charge $O(\sqrt{n})$ of these moves to each vertex of γ.

Otherwise, let σ be an arbitrary useful closed curve chosen according to Lemma 4.4. Suppose the interior tangle of σ has m vertices, s strands, and depth d. Lemma 4.4 implies that $d = O(\sqrt{n})$, and the definition of useful implies that $s \le \sqrt{m}$, which is $O(\sqrt{n})$. Thus, by Corollary 4.3, we can tighten the interior tangle of σ in $O(md + ms) = O(m\sqrt{n})$ moves. This simplification removes at least $m - s^2/2 \ge m/2$ vertices from γ, as the resulting tight tangle has at most $s^2/2$ vertices. Again, for purposes of analysis, we charge $O(\sqrt{n})$ moves to each deleted vertex. We then recursively simplify the resulting closed curve.

In either case, each vertex of γ is charged $O(\sqrt{n})$ moves as it is deleted. Thus, simplification requires at most $O(n^{3/2})$ homotopy moves in total. □

4.4 Efficient Implementation

Here we describe how to implement our curve-simplification algorithm to run in $O(n^{3/2})$ time; in fact, our implementation spends only constant amortized time per homotopy move. We assume that the input curve is given in a data structure that allows fast exploration and modification of plane graphs, such as a quad-edge data structure [31] or a doubly-connected edge list [10]. If the curve is presented as a polygon with m edges, an appropriate graph representation can be constructed in $O(m \log m + n)$ time using classical geometric algorithms [13, 15, 44]; more recent algorithms can be used for piecewise-algebraic curves [20].

Theorem 4.6. *Given a simple closed curve γ in the plane with n vertices, we can compute a sequence of $M = O(n^{3/2})$ homotopy moves that simplifies γ in $O(M)$ time.*

Proof: We begin by labeling each face of γ with its depth, using a breadth-first search of the dual graph in $O(n)$ time. Then we construct the depth contours of γ—the boundaries of the regions $\tilde{R}_j$ from the proof of Lemma 4.4—and organize them into a *contour tree* in $O(n)$ time by brute force. Another $O(n)$-time breadth-first traversal computes the number of strands and the number of interior vertices of every contour's interior tangle; in particular, we identify which depth contours are useful. To complete

the preprocessing phase, we place all the leafmost useful contours into a queue. We can charge the overall $O(n)$ preprocessing time to the $\Omega(n)$ homotopy moves needed to simplify the curve.

As long as the queue of leafmost useful contours is non-empty, we extract one contour σ from this queue and simplify its interior tangle T as follows. Suppose T has m interior vertices.

Following the proof of Theorem 4.5, we first simplify every loop in each strands of T. We identify loops by traversing the strand from one endpoint to the other, marking the vertices as we go; the first time we visit a vertex that has already been marked, we have found a loop α. We can perform each of the homotopy moves required to shrink α in $O(1)$ time, because each such move modifies only a constant-radius boundary of a vertex on α. After the loop is shrunk, we continue walking along the strand starting at the most recently marked vertex.

The second phase of the tangle-simplification algorithm proceeds similarly. We walk around the boundary of T, marking vertices as we go. As soon as we see the second endpoint of any strand γ_i, we pause the walk to straighten γ_i. As before, we can execute each homotopy move used to move γ_i to γ_i' in $O(1)$ time. We then move the boundary of the tangle over the vertices of γ_i', and remove the endpoints of γ_i' from the boundary curve, in $O(1)$ time per vertex.

The only portions of the running time that we have not already charged to homotopy moves are the time spent marking the vertices on each strand and the time to update the tangle boundary after moving a strand aside. Altogether, the uncharged time is $O(m)$, which is less than the number of moves used to tighten T, because the contour σ is useful. Thus, tightening the interior tangle of a useful contour requires $O(1)$ amortized time per homotopy move.

Once the tangle is tight, we must update the queue of useful contours. The original contour σ is still a depth contour in the modified curve, and tightening T only changes the depths of faces that intersect T. Thus, we could update the contour tree in $O(m)$ time, which we could charge to the moves used to tighten T; but in fact, this update is unnecessary, because no contour in the interior of σ is useful. We then walk up the contour tree from σ, updating the number of interior vertices until we find a useful ancestor contour. The total time spent traversing the contour tree for new useful contours is $O(n)$; we can charge this time to the $\Omega(n)$ moves needed to simplify the curve. □

4.5 Multicurves

Finally, we describe how to extend our $O(n^{3/2})$ upper bound to multicurves. Just as in Section 2.3, we distinguish between two variants, depending on whether the target of the simplification is an *arbitrary* set of disjoint cycles or a *particular* set of disjoint cycles. In both cases, our upper bounds match the lower bounds proved in Section 2.3.

First we extend our loop-contraction algorithm from Lemma 4.1 to the multicurve setting. Recall that a *component* of a multicurve γ is any multicurve whose image is a component of the image of γ, and the individual closed curves that comprise γ are its *constituent curves*. The main difficulty is that one component of the multicurve might lie inside a face of another component, making progress on the larger component impossible. To handle this potential obstacle, we simplify the *innermost* components of the multicurve first, and we move isolated simple closed curves toward the outer face as quickly as possible. Figure 4.6 sketches the basic steps of our algorithm when the input multicurve is connected.

Lemma 4.7. *Every n-vertex k-curve γ in the plane can be transformed into k disjoint simple closed curves using at most $3D(\gamma) + 4nk$ homotopy moves.*

Proof: Let γ be an arbitrary k-curve with n vertices. If γ is connected, we either contract and delete a loop, exactly as in Lemma 4.1, or we contract a simple constituent curve to an isolated circle, using essentially the same algorithm. In either case, the number of moves performed is at most $3D(\gamma) - 3D(\gamma')$, where γ' is the multicurve after the contraction. The lemma now follows immediately by induction.

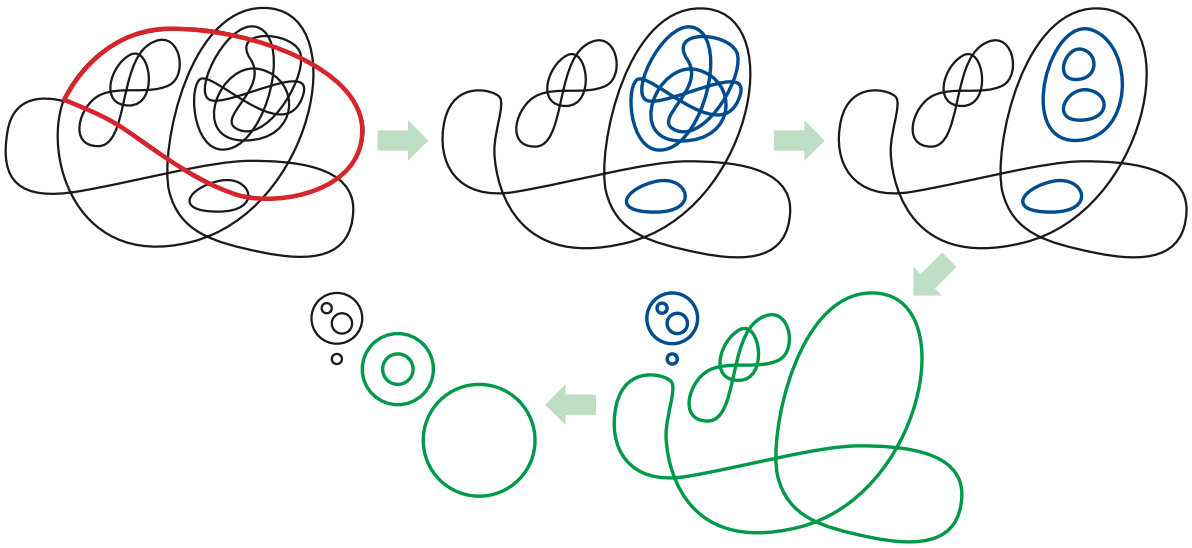

Figure 4.6. Simplifying a connected multicurve: shrink an arbitrary simple loop or cycle, recursively simplify any inner components, translate inner circle clusters to the outer face, and recursively simplify the remaining non-simple components.

We call a component of γ an ***outer component*** if it is incident to the unbounded outer face of γ, and an ***inner component*** otherwise. If γ has more than one outer component, we partition γ into subcurves, each consisting of one outer component γ_o and all inner components located inside faces of γ_o, and we recursively simplify each subcurve independently; the lemma follows by induction. If any outer component is simple, we ignore that component and simplify the rest of γ recursively; again, the lemma follows by induction.

Thus, we can assume without loss of generality that our multicurve γ is disconnected but has only one outer component γ_o, which is non-simple. For each face f of γ_o, let γ_f denote the union of all components inside f. Let n_f and k_f respectively denote the number of vertices and constituent curves of γ_f. Similarly, let n_o and k_o respectively denote the number of vertices and constituent curves of the outer component γ_o.

We first recursively simplify each subcurve γ_f; let κ_f denote the resulting *cluster* of k_f simple closed curves. By the induction hypothesis, this simplification requires at most $3D(\gamma_f) + 4n_f k_f$ homotopy moves. We *translate* each cluster κ_f to the outer face of γ_o by shrinking κ_f to a small ε-ball and then moving the entire cluster along a shortest path in the dual graph of γ_o. This translation requires at most $4n_o k_f$ homotopy moves; each circle in κ_f uses one $2\to0$ move and one $0\to2$ move to cross any edge of γ_o, and in the worst case, the cluster might cross all $2n_0$ edges of γ_o. After all circle clusters are in the outer face, we recursively simplify γ_o using at most $3D(\gamma_o) + 4n_o k_o$ homotopy moves.

The total number of homotopy moves used in this case is

$$\sum_f 3D(\gamma_f) + 3D(\gamma_o) \;+\; \sum_f 4n_f k_f + \sum_f 4n_o k_f + 4n_o k_o.$$

Each face of γ_o has the same depth as the corresponding face of γ, and for each face f of γ_o, each face of the subcurve γ_f has lesser depth than the corresponding face of γ. It follows that

$$\sum_f D(\gamma_f) + D(\gamma_o) \le D(\gamma).$$

Similarly, $\sum_f n_f + n_o = n$ and $\sum_f k_f + k_o = k$. The lemma now follows immediately. □

To reduce the leading term to $O(n^{3/2})$, we extend the definition of a tangle to the intersection of a multicurve γ with a closed disk whose boundary intersects the multicurve transversely away from its vertices, or not at all. Such a tangle can be decomposed into boundary-to-boundary paths, called *open* strands, and closed curves that do not touch the tangle boundary, called *closed* strands. Each closed strand is a constituent curve of γ. A tangle is *tight* if every strand is simple, every pair of open strands intersects at most once, and otherwise all strands are disjoint.

Theorem 4.8. *Every k-curve in the plane with n vertices can be transformed into a set of k disjoint simple closed curves using $O(n^{3/2} + nk)$ homotopy moves.*

Proof: Let γ be an arbitrary k-curve with n vertices. Following the proof of Lemma 4.7, we can assume without loss of generality that γ has a single outer component γ_o, which is non-simple.

When γ is disconnected, we follow the strategy in the previous proof. Let γ_f denote the union of all components inside any face f of γ_o. For each face f, we recursively simplify γ_f and translate the resulting cluster of disjoint circles to the outer face; when all faces are empty, we recursively simplify γ_o. The theorem now follows by induction.

When γ is non-simple and connected, we follow the useful closed curve strategy from Theorem 4.5. We define a closed curve σ to be useful for γ if the interior tangle of σ has its number of vertices at least the square of the number of *open* strands; then the proof of Lemma 4.4 applies to connected multicurves with no modifications. So let T be a tangle with m vertices, $s \le \sqrt{m}$ open strands, ℓ closed strands, and depth $d = O(\sqrt{n})$. We straighten T in two phases, almost exactly as in Section 4.2, contracting loops and simple closed strands in the first phase, and straightening open strands in the second phase.

In the first phase, contracting one loop or simple closed strand uses at most $3D(T) - 3D(T')$ homotopy moves, where T' is the tangle after contraction. After each contraction, if T' is disconnected—in particular, if we just contracted a simple closed strand—we simplify and extract any isolated components as follows. Let T'_o denote the component of T' that includes the boundary cycle, and for each face f of T'_o, let γ_f denote the union of all components of T' inside f. We simplify each multicurve γ_f using the algorithm from Lemma 4.7—*not recursively!*—and then translate the resulting cluster of disjoint circles *to the outer face of* γ. See Figure 4.7. Altogether, simplifying and translating these subcurves requires at most $3D(T') - 3D(T'') + 4n\sum_f k_f$ homotopy moves, where T'' is the resulting tangle.

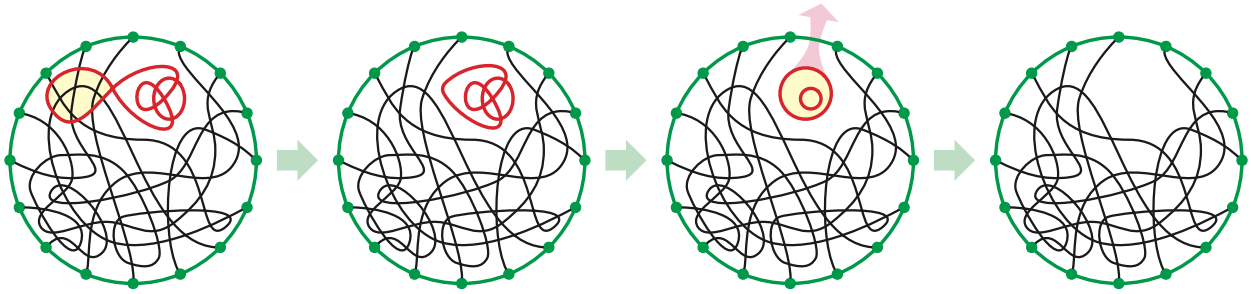

Figure 4.7. Whenever shrinking a loop or simple closed strand disconnects the tangle, simplify each isolated component and translate the resulting cluster of circles to the outer face of the entire multicurve.

The total number of moves performed in the first phase is at most $3D(T) + 4m\ell = O(m\sqrt{n} + n\ell)$. The first phase ends when the tangle consists entirely of simple open strands. Thus, the second phase straightens the remaining open strands exactly as in the proof of Lemma 4.2; the total number of moves in this phase is $O(ms) = O(m\sqrt{n})$. We charge $O(\sqrt{n})$ time to each deleted vertex and $O(n)$ time to each constituent curve that was simplified and translated outward. We then recursively simplify the remaining multicurve, ignoring any outer circle clusters.

Altogether, each vertex of γ is charged $O(\sqrt{n})$ time as it is deleted, and each constituent curve of γ is charged $O(n)$ time as it is translated outward. □

With $O(k^2)$ additional homotopy moves, we can transform the resulting set of k disjoint circles into k nested or unnested circles.

Theorem 4.9. *Any k-curve with n vertices in the plane can be transformed into k nested (or unnested) simple closed curves using $O(n^{3/2} + nk + k^2)$ homotopy moves.*

Corollary 4.10. *Any k-curve with at most n vertices in the plane can be transformed into any other k-curve with at most n vertices using $O(n^{3/2} + nk + k^2)$ homotopy moves.*

Theorems 2.6 and 2.8 and Corollary 2.9 imply that these upper bounds are tight in the worst case for all possible values of n and k. As in the lower bounds, the $O(k^2)$ terms are redundant for connected multicurves.

More careful analysis implies that any k-curve with n vertices and depth d can be simplified in $O(n \min\{d, n^{1/2}\} + k \min\{d, n\})$ homotopy moves, or transformed into k unnested circles using $O(n \min\{d, n^{1/2}\} + k \min\{d, n\} + k \min\{d, k\})$ homotopy moves. Moreover, these upper bounds are tight for all possible values of n, k, and d. We leave the details of this extension as an exercise for the reader.

5 Higher-Genus Surfaces

Finally, we consider the natural generalization of our problem to closed curves on orientable surfaces of higher genus. Because these surfaces have non-trivial topology, not every closed curve is homotopic to a single point or even to a simple curve. A closed curve is ***contractible*** if it is homotopic to a single point. We call a closed curve ***tight*** if it has the minimum number of self-intersections in its homotopy class.

5.1 Lower Bounds

Although defect was originally defined as an invariant of *planar* curves, Polyak's formula $\mathit{defect}(\gamma) = -2\sum_{x \lozenge y} \operatorname{sgn}(x)\operatorname{sgn}(y)$ extends naturally to closed curves on any orientable surface; homotopy moves change the resulting invariant exactly as described in Figure 2.1. Thus, Lemma 2.1 immediately generalizes to any orientable surface as follows.

Lemma 5.1. *Let γ and γ' be arbitrary closed curves that are homotopic on an arbitrary orientable surface. Transforming γ into γ' requires at least $|\mathit{defect}(\gamma) - \mathit{defect}(\gamma')|/2$ homotopy moves.*

The following construction implies a quadratic lower bound for tightening noncontractible curves on orientable surfaces with any positive genus.

Lemma 5.2. *For any positive integer n, there is a closed curve on the torus with n vertices and defect $\Omega(n^2)$ that is homotopic to a simple closed curve but not contractible.*

Proof: Without loss of generality, suppose n is a multiple of 8. The curve γ is illustrated on the left in Figure 5.1. The torus is represented by a rectangle with opposite edges identified. We label three points a, b, c on the vertical edge of the rectangle and decompose the curve into a red path from a to b, a green path from b to c, and a blue path from c to a. The red and blue paths each wind vertically around the torus, first $n/8$ times in one direction, and then $n/8$ times in the opposite direction.

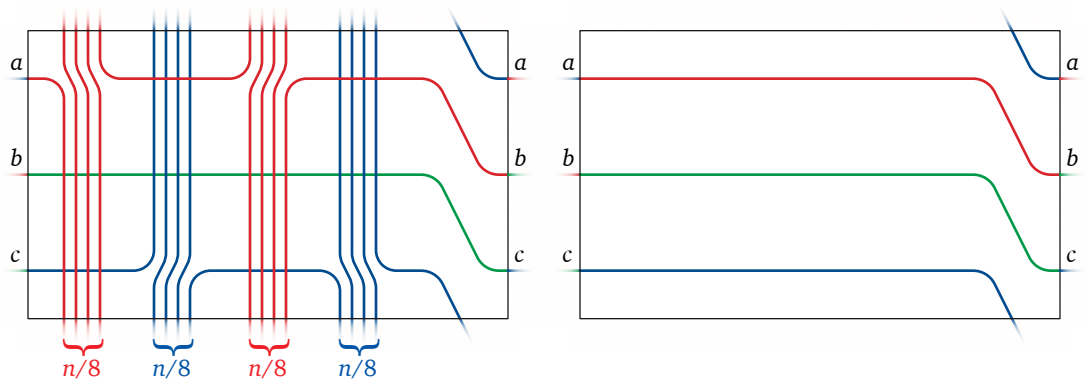

Figure 5.1. A curve γ on the torus with defect $\Omega(n^2)$ and a simple curve homotopic to γ.

As in previous proofs, we compute the defect of γ by describing a sequence of homotopy moves that simplifies the curve, while carefully tracking the changes in the defect that these moves incur. We can unwind one turn of the red path by performing one $2\to0$ move, followed by $n/8$ $3\to3$ moves, followed by one $2\to0$ move, as illustrated in Figure 5.2. Repeating this sequence of homotopy moves $n/8$ times removes all intersections between the red and green paths, after which a sequence of $n/4$ $2\to0$ moves straightens the blue path, yielding the simple curve shown on the right in Figure 5.1. Altogether, we perform $n^2/64+n/2$ homotopy moves, where each $3\to3$ move increases the defect of the curve by 2 and each $2\to0$ move decreases the defect of the curve by 2. We conclude that $\mathit{defect}(\gamma) = -n^2/32+n$. □

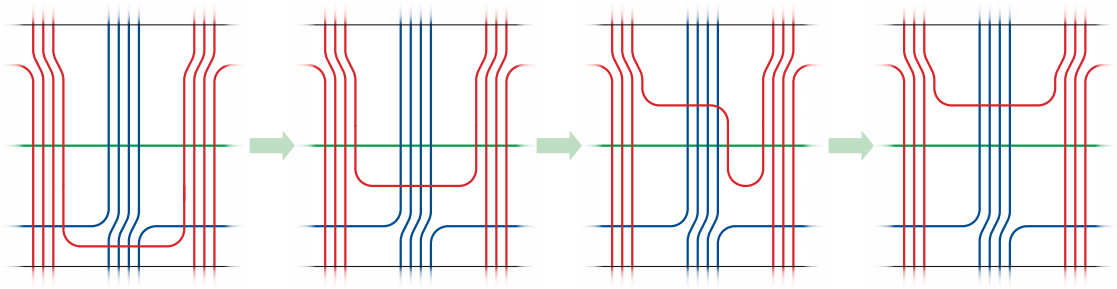

Figure 5.2. Unwinding one turn of the red path.

Theorem 5.3. *Tightening a closed curve with n crossings on a torus requires $\Omega(n^2)$ homotopy moves in the worst case, even if the curve is homotopic to a simple curve.*

5.2 Upper Bounds

Hass and Scott proved that any non-simple closed curve on any orientable surface that is homotopic to a simple closed curve contains either a simple (in fact empty) contractible loop or a simple contractible bigon [33, Theorem 1]. It follows immediately that any such curve can be simplified in $O(n^2)$ moves using Steinitz's algorithm; Theorem 5.3 implies that the upper bound is tight for non-contractible curves.

For the most general setting, where the given curve is not necessarily homotopic to a simple closed curve, we are not even aware of a *polynomial* upper bound! Steinitz's algorithm does not work here; there are curves with excess self-intersections but no simple contractible loops or bigons [33]. Hass and Scott [34] and De Graff and Schrijver [28] independently proved that any closed curve on any

surface can be simplified using a *finite* number of homotopy moves that never increase the number of self-intersections. Both proofs use discrete variants of curve-shortening flow; for sufficiently well-behaved curves and surfaces, results of Grayson [29] and Angenent [6] imply a similar result for differential curvature flow. Unfortunately, without further assumptions about the precise geometries of both the curve and the underlying surface, the number of homotopy moves cannot be bounded by any function of the number of crossings; even in the plane, there are closed curves with three crossings for which curve-shortening flow alternates between a 3→3 move and its inverse arbitrarily many times. Paterson [48] describes a combinatorial algorithm to compute a tightening sequence of homotopy moves without such reversals, but she offers no analysis of her algorithm.

We conjecture that any arbitrary curves (or even multicurves) on any surface can be simplified with at most $O(n^2)$ homotopy moves.

Acknowledgements. We would like to thank Nathan Dunfield, Joel Hass, Bojan Mohar, and Bob Tarjan for encouragement and helpful discussions, and Joe O'Rourke for asking about multiple curves. We would also like to thank the anonymous referees for several suggestions that improved the paper.

References

[1] Francesca Aicardi. Tree-like curves. *Singularities and Bifurcations*, 1–31, 1994. Advances in Soviet Mathematics 21, Amer. Math. Soc.

[2] Sheldon B. Akers, Jr. The use of wye-delta transformations in network simplification. *Oper. Res.* 8(3):311–323, 1960.

[3] James W. Alexander. Combinatorial analysis situs. *Trans. Amer. Math. Soc.* 28(2):301–326, 1926.

[4] James W. Alexander and G. B. Briggs. On types of knotted curves. *Ann. Math.* 28(1/4):562–586, 1926–1927.

[5] Sarah R. Allen, Luis Barba, John Iacono, and Stefan Langerman. Incremental Voronoi diagrams. *Proc. 32nd Int. Symp. Comput. Geom.*, 15:1–15:16, 2016. Leibniz International Proceedings in Informatics 51. ⟨http://drops.dagstuhl.de/opus/volltexte/2016/5907⟩. arXiv:1603.08485.

[6] Sigurd Angenent. Parabolic equations for curves on surfaces: Part II. Intersections, blow-up and generalized solutions. *Ann. Math.* 133(1):171–215, 1991.

[7] Dan Archdeacon, Charles J. Colbourn, Isidoro Gitler, and J. Scott Provan. Four-terminal reducibility and projective-planar wye-delta-wye-reducible graphs. *J. Graph Theory* 33(2):83–93, 2000.

[8] Vladimir I. Arnold. Plane curves, their invariants, perestroikas and classifications. *Singularities and Bifurcations*, 33–91, 1994. Adv. Soviet Math. 21, Amer. Math. Soc.

[9] Vladimir I. Arnold. *Topological Invariants of Plane Curves and Caustics*. University Lecture Series 5. Amer. Math. Soc., 1994.

[10] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*, 3rd edition. Springer-Verlag, 2008.

[11] Hsien-Chih Chang and Jeff Erickson. Electrical reduction, homotopy moves, and defect. Preprint, October 2015. arXiv:1510.00571.

[12] Hsien-Chih Chang and Jeff Erickson. Untangling planar curves. *Proc. 32nd Int. Symp. Comput. Geom.*, 29:1–29:15, 2016. Leibniz International Proceedings in Informatics 51. ⟨http://drops.dagstuhl.de/opus/volltexte/2016/5921⟩.

[13] Bernard Chazelle and Herbert Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *J. ACM* 39(1):1–54, 1992.

[14] Sergei Chmutov, Sergei Duzhin, and Jacob Mostovoy. *Introduction to Vassiliev knot invariants*. Cambridge Univ. Press, 2012. ⟨http://www.pdmi.ras.ru/~duzhin/papers/cdbook⟩. arXiv:1103.5628.

[15] Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.* 4:387–421, 1989.

[16] Charles J. Colbourn, J. Scott Provan, and Dirk Vertigan. A new approach to solving three combinatorial enumeration problems on planar graphs. *Discrete Appl. Math.* 60:119–129, 1995.

[17] Yves Colin de Verdière, Isidoro Gitler, and Dirk Vertigan. Réseaux électriques planaires II. *Comment. Math. Helvetici* 71:144–167, 1996.

[18] John H. Conway. An enumeration of knots and links, and some of their algebraic properties. *Computational Problems in Abstract Algebra*, 329–358, 1970. Pergamon Press.

[19] Alexander Coward and Marc Lackenby. An upper bound on Reidemeister moves. *Amer. J. Math.* 136(4):1023–1066, 2014. arXiv:1104.1882.

[20] Arno Eigenwillig and Michael Kerber. Exact and efficient 2D-arrangements of arbitrary algebraic curves. *Proc. 19th Ann. ACM-SIAM Symp. Discrete Algorithms*, 122–131, 2008.

[21] G. V. Epifanov. Reduction of a plane graph to an edge by a star-triangle transformation. *Dokl. Akad. Nauk SSSR* 166:19–22, 1966. In Russian. English translation in *Soviet Math. Dokl.* 7:13–17, 1966.

[22] Chaim Even-Zohar, Joel Hass, Nati Linial, and Tahl Nowik. Invariants of random knots and links. *Discrete & Computational Geometry* 56(2):274–314, 2016. arXiv:1411.3308.

[23] Thomas A. Feo. *I. A Lagrangian Relaxation Method for Testing The Infeasibility of Certain VLSI Routing Problems. II. Efficient Reduction of Planar Networks For Solving Certain Combinatorial Problems*. Ph.D. thesis, Univ. California Berkeley, 1985. ⟨http://search.proquest.com/docview/303364161⟩.

[24] Thomas A. Feo and J. Scott Provan. Delta-wye transformations and the efficient reduction of two-terminal planar graphs. *Oper. Res.* 41(3):572–582, 1993.

[25] George K. Francis. The folded ribbon theorem: A contribution to the study of immersed circles. *Trans. Amer. Math. Soc.* 141:271–303, 1969.

[26] Carl Friedrich Gauß. Nachlass. I. Zur Geometria situs. *Werke*, vol. 8, 271–281, 1900. Teubner. Originally written between 1823 and 1840.

[27] Isidoro Gitler. *Delta-wye-delta Transformations: Algorithms and Applications*. Ph.D. thesis, Department of Combinatorics and Optimization, University of Waterloo, 1991.

[28] Maurits de Graaf and Alexander Schrijver. Making curves minimally crossing by Reidemeister moves. *J. Comb. Theory Ser. B* 70(1):134–156, 1997.

[29] Matthew A. Grayson. Shortening embedded curves. *Ann. Math.* 129(1):71–111, 1989.

[30] Branko Grünbaum. *Convex Polytopes*. Monographs in Pure and Applied Mathematics XVI. John Wiley & Sons, 1967.

[31] Leonidas J. Guibas and Jorge Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graphics* 4(2):75–123, 1985.

[32] Joel Hass and Tal Nowik. Unknot diagrams requiring a quadratic number of Reidemeister moves to untangle. *Discrete Comput. Geom.* 44(1):91–95, 2010.

[33] Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel J. Math.* 51:90–120, 1985.

[34] Joel Hass and Peter Scott. Shortening curves on surfaces. *Topology* 33(1):25–43, 1994.

[35] Chuichiro Hayashi and Miwa Hayashi. Minimal sequences of Reidemeister moves on diagrams of torus knots. *Proc. Amer. Math. Soc.* 139:2605–2614, 2011. arXiv:1003.1349.

[36] Chuichiro Hayashi, Miwa Hayashi, and Tahl Nowik. Unknotting number and number of Reidemeister moves needed for unlinking. *Topology Appl.* 159:1467–1474, 2012. arXiv:1012.4131.

[37] Chuichiro Hayashi, Miwa Hayashi, Minori Sawada, and Sayaka Yamada. Minimal unknotting sequences of Reidemeister moves containing unmatched RII moves. *J. Knot Theory Ramif.* 21(10):1250099 (13 pages), 2012. arXiv:1011.3963.

[38] Heinz Hopf. Über die Drehung der Tangenten und Sehnen ebener Kurven. *Compositio Math.* 2:50–62, 1935.

[39] Noburo Ito and Yusuke Takimura. (1,2) and weak (1,3) homotopies on knot projections. *J. Knot Theory Ramif.* 22(14):1350085 (14 pages), 2013. Addendum in *J. Knot Theory Ramif.* 23(8):1491001 (2 pages), 2014.

[40] Arthur Edwin Kennelly. Equivalence of triangles and three-pointed stars in conducting networks. *Electrical World and Engineer* 34(12):413–414, 1899.

[41] Mikhail Khovanov. Doodle groups. *Trans. Amer. Math. Soc.* 349(6):2297–2315, 1997.

[42] Marc Lackenby. A polynomial upper bound on Reidemeister moves. *Ann. Math.* 182(2):491–564, 2015. arXiv:1302.0180.

[43] Alfred Lehman. Wye-delta transformations in probabilistic network. *J. Soc. Indust. Appl. Math.* 11:773–805, 1963.

[44] Ketan Mulmuley. A fast planar partition algorithm, I. *J. Symbolic Comput.* 10(3–4):253–280, 1990.

[45] Hiroyuki Nakahara and Hiromitsu Takahashi. An algorithm for the solution of a linear system by Δ-Y transformations. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences* E79-A(7):1079–1088, 1996. Special Section on Multi-dimensional Mobile Information Network.

[46] Steven D. Noble and Dominic J. A. Welsh. Knot graphs. *J. Graph Theory* 34(1):100–111, 2000.

[47] Tahl Nowik. Complexity of planar and spherical curves. *Duke J. Math.* 148(1):107–118, 2009.

[48] Jane M. Paterson. A combinatorial algorithm for immersed loops in surfaces. *Topology Appl.* 123:205–234, 2002.

[49] Michael Polyak. Invariants of curves and fronts via Gauss diagrams. *Topology* 37(5):989–1009, 1998.

[50] Kurt Reidemeister. Elementare Begründung der Knotentheorie. *Abh. Math. Sem. Hamburg* 5:24–32, 1927.

[51] Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Comb. Theory Ser. B* 52(2):153–190, 1991.

[52] Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *J. Comb. Theory Ser. B* 62(2):232–348, 1994.

[53] Alexander Russell. The method of duality. *A Treatise on the Theory of Alternating Currents*, chapter XVII, 380–399, 1904. Cambridge Univ. Press.

[54] Xiaohuan Song. Implementation issues for Feo and Provan's delta-wye-delta reduction algorithm. M.Sc. Thesis, University of Victoria, 2001.

[55] Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen* III.AB(12):1–139, 1916.

[56] Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder: unter Einschluß der Elemente der Topologie*. Grundlehren der mathematischen Wissenschaften 41. Springer-Verlag, 1934. Reprinted 1976.

[57] James A. Storer. On minimal node-cost planar embeddings. *Networks* 14(2):181–212, 1984.

[58] Klaus Truemper. On the delta-wye reduction for planar graphs. *J. Graph Theory* 13(2):141–148, 1989.

[59] Klaus Truemper. *Matroid Decomposition*. Academic Press, 1992.

[60] Leslie S. Valiant. Universality considerations in VLSI circuits. *IEEE Trans. Comput.* C-30(2):135–140, 1981.

[61] Gert Vegter. Kink-free deformation of polygons. *Proceedings of the 5th Annual Symposium on Computational Geometry*, 61–68, 1989.